

# Robust Promptable Video Object Segmentation

Sohyun Lee<sup>1</sup> Yeho Gwon<sup>1</sup> Lukas Hoyer<sup>2</sup> Konrad Schindler<sup>3</sup> Christos Sakaridis<sup>3</sup> Suha Kwak<sup>1</sup>

<sup>1</sup>POSTECH

<sup>2</sup>Google

<sup>3</sup>ETH Zürich

[https://sohyun-1.github.io/RobustPVOS\\_project\\_page/](https://sohyun-1.github.io/RobustPVOS_project_page/)

## Abstract

The performance of promptable video object segmentation (PVOS) models substantially degrades under input corruptions, which prevents PVOS deployment in safety-critical domains. This paper offers the first comprehensive study on robust PVOS (RobustPVOS). We first construct a new, comprehensive benchmark with two real-world evaluation datasets of 351 video clips and more than 2,500 object masks under real-world adverse conditions. At the same time, we generate synthetic training data by applying diverse and temporally varying corruptions to existing VOS datasets. Moreover, we present a new RobustPVOS method, dubbed Memory-object-conditioned Gated-rank Adaptation (MoGA). The key to successfully performing RobustPVOS is two-fold: effectively handling object-specific degradation and ensuring temporal consistency in predictions. MoGA leverages object-specific representations maintained in memory across frames to condition the robustification process, which allows the model to handle each tracked object differently in a temporally consistent way. Extensive experiments on our benchmark validate MoGA’s efficacy, showing consistent and significant improvements across diverse corruption types on both synthetic and real-world datasets, establishing a strong baseline for future RobustPVOS research. Our benchmark is publicly available at [our project page](#).

## 1. Introduction

Promptable video object segmentation (PVOS) has emerged as a powerful paradigm, enabling users to segment and track arbitrary objects in videos given free-form prompts [8, 28, 33, 37, 46]. The recent introduction of SAM2 [37] marks a significant breakthrough in this field, achieving impressive zero-shot performance by extending the success of promptable segmentation from images to videos. However, we empirically find that the performance of SAM2 substantially degrades under input corruptions caused by noise, blur, low illumination, and adverse weather. This lack of robust-




Figure 1. Overview of (a) RobustPVOS and (b) our benchmark. RobustPVOS is the task of tracking and segmenting objects, indicated by initial prompts, across frames despite adverse conditions.

ness poses critical challenges for deploying PVOS in safety-critical domains such as autonomous vehicles and robotics, where adverse conditions are inevitable, but has not been adequately studied or addressed to date.

Motivated by this, we introduce the first comprehensive study on robust promptable video object segmentation (RobustPVOS), whose goal and benchmark composition are illustrated in Figure 1. Our primary contribution is establishing RobustPVOS as a well-defined research direction through a carefully designed benchmark that enables development and systematic evaluation of RobustPVOS models. We construct two real-world evaluation datasets comprising 351 video clips with over 2,500 annotated object masks from existing video collections captured under natural adverse conditions including rain, fog, snow, and nighttime scenarios [16, 38, 39]. To our knowledge, this is the first effort to provide dense, object-level annotations for promptable video segmentation under real-world corruptions. Additionally, we create synthetic training data by applying eight diverse corruption types with temporal variations [4, 13, 19] to annotated videos from established VOS datasets [7, 34, 47], enabling model training and evaluation under controlled conditions. These datasets provide a valuable testbed for assessing model performance across diverse degradation scenarios, from controlled synthetic corruptions to naturally occurring adverse conditions.

Beyond the benchmark, we further investigate method-

ological approaches for RobustPVOS. One potential approach to RobustPVOS is applying existing robust image segmentation methods [3, 4, 24, 26, 41] to individual video frames. While these methods effectively handle corrupted images, they face fundamental limitations in PVOS, since real-world videos exhibit complex degradation patterns that vary both spatially and temporally: different objects within the same scene are affected differently (*e.g.*, in fog, distant objects are barely visible while close objects remain clearly visible), and degradation patterns change across frames. Robust image segmentation models that process frames independently often yield temporally inconsistent segmentation results. Also, since they process the entire area of each frame uniformly for improving robustness, they fail to consider the varying effects of corruptions on different objects within the same scene.

To address these challenges, we propose Memory-object-conditioned Gated-rank Adaptation, dubbed MoGA, as an effective baseline method for RobustPVOS. Our key insight is that object-specific representations maintained across frames, which are inherent in modern video segmentation models [33, 37, 49], naturally capture how each object is uniquely affected by degradations over time. By conditioning the robustification process on the representations of these objects in the memory, MoGA enables the model to handle each tracked object differently while maintaining temporal consistency. This design demonstrates that memory-based conditioning can substantially improve robustness in PVOS.

Extensive experiments on our benchmark validate the effectiveness of our approach. MoGA achieves consistent improvements across diverse corruption types on both synthetic and real-world datasets, while maintaining competitive performance on clear videos. Our benchmark and baseline together establish foundations for future research in RobustPVOS. In summary, our contribution is three-fold:

- We initiate the study of RobustPVOS, establishing it as a critical research direction for real-world video understanding applications.
- We construct a comprehensive RobustPVOS benchmark, including real-world evaluation datasets with object-level annotations under adverse conditions and synthetic training data with controlled temporal degradations.
- We present MoGA as an effective baseline method that addresses the unique challenges of RobustPVOS through object-specific and temporally consistent robustification.

## 2. Related Work

### 2.1. Video Object Segmentation

Video object segmentation (VOS) aims to identify and track target objects throughout video sequences [37]. Traditional methods rely on dense optical flow [10, 42–44] to maintain

temporal consistency. Recent approaches leverage memory modules to store object representations over videos [33, 37, 40, 49], enabling stable tracking. RMem [49] investigates the effect of a restricted memory bank, while SAM-I2V [33] focuses on exploiting pre-trained SAM by introducing a memory prompt generator. Notably, Cutie [5] employs object queries to construct an object-level memory in a top-down manner. As PVOS has emerged as a new paradigm for prompt-based object segmentation [12, 33, 37], UniVS [28] accepts visual and text prompts and segments objects by averaging their features. However, these methods are primarily designed for clear videos and struggle under adverse conditions [29], motivating our work on RobustPVOS.

### 2.2. Promptable Segmentation

The Segment Anything Model (SAM) [20] revolutionized image segmentation by enabling interactive mask generation. SAM accepts diverse prompt types such as points, boxes, and masks, demonstrating remarkable zero-shot generalization across domains. HQ-SAM [17] extends SAM for high-quality segmentation by introducing a learnable output token with intricate structures. SAM2 [37] extends the promptable segmentation ability of SAM to videos by introducing a streaming memory architecture that propagates the masks of prompted objects across frames. SEEM [50] and COSINE [30] support diverse prompt types, including text and images, through a unified decoding mechanism with learnable memory prompts. CLIPSeg [32] pioneers text-based segmentation by combining CLIP with a transformer decoder for both text queries and visual examples, while REM [2] adopts a text-to-image generative diffusion model to segment a moving object in a video. While these methods excel on clear data, their robustness under corruptions remains unexplored, particularly for video applications where temporal consistency is crucial.

### 2.3. Robustness

Robustness to input corruptions has been extensively studied for the image domain [3, 4, 20–25]. For semantic segmentation, recent works propose various strategies, such as FIFO [21], which learns fog-invariant features, and FreD [3], which employs frequency-domain analysis. RobustSAM [4] improves the robustness of SAM [20] by introducing anti-degradation modules. GaRA-SAM [24] introduces gated-rank adaptation, achieving input-adaptive robustness for SAM through selective (rank-1) component activation. Universal restoration methods [1, 26, 35, 41] restore corrupted images and show that restoration improves robust recognition. However, they process frames independently, so they cannot model temporal consistency in videos. For robustness in the video domain, VPSeg [11] exploits vanishing point priors for robust video semantic segmentation in driving scenes. Event cameras are also used

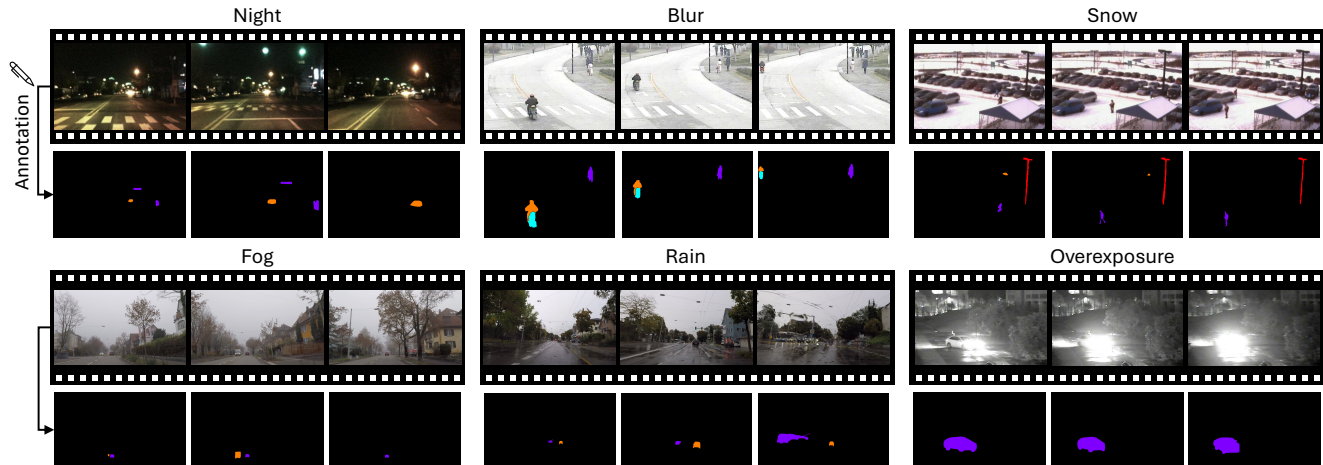


Figure 2. Example images and annotated object masks of the real-world evaluation dataset.

Table 1. Statistics of our benchmark for RobustPVOS. The top block lists real-world evaluation sets; the bottom blocks summarize the synthetic training set and the synthetic evaluation sets.

Dataset	Clips	Frames	Objects
<b>Real-world evaluation set</b>			
ACDC-Video	149	3,259	613
MVSeg	202	13,581	1,930
<b>Synthetic training set</b>			
MOSE-C + DAVIS-C + YouTube-VOS-C	46,768	1,774,560	3,872,048
<b>Synthetic evaluation set</b>			
YouTube-VOS-C	507	13,710	25,574

to better understand motion in videos [18, 27], particularly in challenging conditions such as motion blur and low illumination. Although they handle the temporal dynamics of videos, they still fall short in RobustPVOS due to its multi-object nature. Our work fills this gap by introducing memory-based conditioning, specifically designed for temporally consistent and object-specific robustification.

### 3. RobustPVOS Benchmark

We present the first RobustPVOS benchmark suite that includes (1) manually annotated real-world evaluation datasets under adverse conditions and (2) a synthetic corruption dataset with temporally varying degradation patterns. The overall composition of our benchmark is summarized in Table 1, and qualitative examples of the real-world evaluation dataset are presented in Figure 2.

**Real-World Evaluation Datasets.** We curate two real-world PVOS test sets by collecting videos captured under natural adverse conditions from the ACDC-Video [39] and MVSeg [16] datasets. While these datasets provide masks, they were not originally designed for PVOS: MVSeg provides only class-level masks, and ACDC offers instance-level masks without temporally consistent tracking. We therefore annotate dense, per-object, pixel-level masks with consistent instance identities across frames, labeling instance IDs and correcting low-quality masks. We also manually excluded clips with mild adverse conditions or static objects to ensure suitability for RobustPVOS evaluation. As shown in Table 1, we collect 149 clips (3,259 frames) captured in fog, snow, rain, and nighttime conditions from ACDC-Video [39], annotating 613 object masks on key frames. From MVSeg [16], we annotate 202 video clips (13,581 frames) across various conditions and release them as a public PVOS dataset, providing 1,930 annotated object masks on key frames (Table 1); for RobustPVOS evaluation, we further construct a dedicated test subset, dubbed MVSeg-adv, by retaining only the clips captured under challenging conditions such as low light, rain, snow, noise, and motion blur, excluding those with mild adverse conditions or static objects.

**Synthetic Corruption Dataset.** For both training and controlled evaluation, we further construct a synthetic dataset by applying eight types of corruption (*i.e.*, color jitter, Gaussian noise, ISO noise, motion blur, resampling blur, fog, rain, and snow) to existing VOS datasets. Using corruption synthesis algorithms [4, 13] with Fourier-based temporal modulation [19], we vary corruption intensity smoothly across frames to simulate degradation patterns in real videos. For training, we corrupt videos from MOSE [7], YouTube-VOS [47], and DAVIS [34], gener-

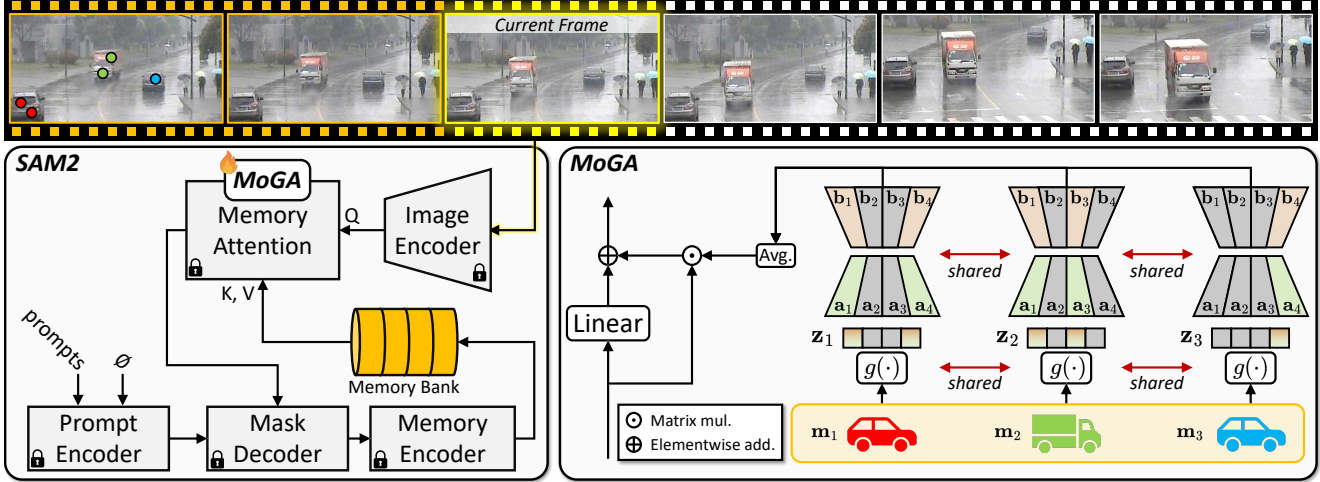


Figure 3. Overview of MoGA integrated into SAM2 [37]. *Top*: an example input video under adverse weather conditions. Frames with orange outlines are previously processed and stored in the memory bank, while the yellow outline indicates the current frame. *Bottom left*: MoGA modules are integrated into the memory attention mechanism of SAM2. *Bottom right*: our memory-conditioned gating mechanism where object pointers from the memory bank guide the selection of (rank-1) components of the shared adapter, enabling object-specific robustness while maintaining temporal consistency. Best viewed in color.

ating 46,768 clips (1,774,560 frames) with 3,872,048 object masks. For evaluation, we create YouTube-VOS-C, a test set containing 507 corrupted clips (13,710 frames) with 25,574 object masks from the YouTube-VOS validation set.

**Benchmarking Protocol.** For each benchmark video, we provide prompts for all annotated objects in the first frame. Models must segment and track these objects throughout the entire sequence. Their performance is evaluated using the standard VOS metrics [37, 47]: region similarity ( $\mathcal{J}$ ), contour accuracy ( $\mathcal{F}$ ), and their combination ( $\mathcal{J}\&\mathcal{F}$ ).

## 4. Method

We present Memory-object-conditioned Gated-rank Adaptation (MoGA) for RobustPVOS. The overall architecture of MoGA is illustrated in Figure 3. The remainder of this section provides preliminaries on low-rank adaptation in Section 4.1 and elaborates on MoGA in Section 4.2.

### 4.1. Preliminaries on Low-Rank Adaptation

Low-rank Adaptation (LoRA) [14] is a parameter-efficient fine-tuning technique that freezes a pre-trained weight matrix  $\mathbf{W}_0 \in \mathbb{R}^{D \times K}$  and introduces a learnable low-rank adapter  $\Delta\mathbf{W} = \mathbf{B}\mathbf{A}$ , where  $\mathbf{B} \in \mathbb{R}^{D \times R}$  and  $\mathbf{A} \in \mathbb{R}^{R \times K}$  with rank  $R \ll \min(D, K)$ . This adapter modifies the forward pass as:

$$\mathbf{h} = \mathbf{W}_0\mathbf{x} + \mathbf{B}\mathbf{A}\mathbf{x}, \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^K$  is the input and  $\mathbf{h} \in \mathbb{R}^D$  is the output. Recent work [24] has shown that decomposing a low-rank

adapter into components and selectively activating them can improve robustness to input degradation, but it does not guarantee object-wise temporally consistent robustification.

### 4.2. MoGA

For RobustPVOS, MoGA leverages object representations maintained across frames in modern video segmentation architectures [33, 37, 49]. These architectures employ a memory bank as a temporal storage that accumulates object features from processed frames to enable consistent tracking. MoGA decomposes the weight matrix of the low-rank adapter, *i.e.*,  $\Delta\mathbf{W}$  of Section 4.1, into (rank-1) components, and selectively activates them using object-specific representations from the memory bank that stores information of the object from previous frames. This design enables the adapter to handle distinct objects differently while ensuring the temporal consistency.

**(Rank-1) Component Decomposition.** Following previous work [24], we decompose the learnable weight matrix  $\Delta\mathbf{W} = \mathbf{B}\mathbf{A}$  into  $R$  (rank-1) components  $\{\mathbf{a}_i, \mathbf{b}_i\}_{i=1}^R$ , where  $\mathbf{a}_i \in \mathbb{R}^K$ ,  $\mathbf{b}_i \in \mathbb{R}^D$ , and  $\Delta\mathbf{W} = \sum_{i=1}^R \mathbf{b}_i\mathbf{a}_i^\top$ . This decomposition enables a flexible construction of the weight matrix based on the input characteristics.

**Memory-object-conditioned Gating.** The key innovation in MoGA is conditioning the gating mechanism on object information stored in the memory bank. SAM2 [37] maintains object pointers  $\mathcal{M} = \{\mathbf{m}_o\}_{o=1}^O$  in the memory bank which accumulate information across frames. Each

Table 2. Evaluation results on real-world and synthetic datasets. We evaluate on MVSeg, ACDC-Video, and YouTube-VOS-C and its counterpart.

Method	MVSeg-adv [16]			ACDC-Video [39]			YouTube-VOS-C			YouTube-VOS [47]		
	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
SAM2 [37]	69.6	60.5	78.6	63.5	49.8	77.2	78.7	77.2	80.2	82.2	80.4	83.9
URIE [41]+SAM2	69.6	60.7	78.4	60.9	46.4	75.3	78.6	77.1	80.2	-	-	-
AirNet [26]+SAM2	69.1	60.4	77.7	59.8	46.2	73.5	78.6	77.1	80.0	-	-	-
GaRA [24]+SAM2	69.7	61.4	78.0	61.3	47.4	75.2	78.1	76.3	79.9	79.8	78.1	81.4
MoGA+SAM2	<b>71.8</b>	<b>62.9</b>	<b>80.7</b>	<b>64.5</b>	<b>50.3</b>	<b>78.8</b>	<b>79.9</b>	<b>78.3</b>	<b>81.5</b>	<b>82.6</b>	<b>80.7</b>	<b>84.4</b>

object pointer  $\mathbf{m}_o \in \mathbb{R}^d$  encodes the historical characteristics specific to object  $o$ . We introduce a gating module  $g(\cdot)$  that takes an object pointer as input and computes object-specific binary gating masks  $\mathbf{z}_o = g(\mathbf{m}_o) \in \{0, 1\}^R$ . The gating module, which consists of a three-layer MLP followed by Gumbel-Sigmoid relaxation, is shared but applied per-object using its corresponding object pointer  $\mathbf{m}_o$ :

$$\alpha_o = \text{MLP}(\mathbf{m}_o) \quad (2)$$

To enable differentiable learning with discrete selection, we apply Gumbel-sigmoid relaxation [15] to the logit vectors:

$$\tilde{z}_{o,i} = \sigma \left( \frac{1}{\tau} (\alpha_{o,i} + G_i) \right), \quad i = 1, \dots, R, \quad (3)$$

where  $G_i \sim \text{Gumbel}(0, 1)$ ,  $\tau$  is the temperature parameter, and  $\sigma$  is the sigmoid function. During training, we use hard thresholding in the forward pass while maintaining differentiability through the straight-through estimator:

$$z_{o,i} = \begin{cases} \mathbb{I}[\tilde{z}_{o,i} > 0.5] & \text{(forward)} \\ \tilde{z}_{o,i} & \text{(backward)} \end{cases} \quad (4)$$

At inference, gating becomes deterministic, with  $z_{o,i} = \mathbb{I}[\sigma(\alpha_{o,i}) > 0.5]$ . MoGA computes the output  $\mathbf{h}$  as

$$\mathbf{h} = \mathbf{W}_0 \mathbf{x} + \frac{1}{O} \sum_{o=1}^O \Delta \mathbf{W}_o \mathbf{x} \quad (5)$$

$$= \mathbf{W}_0 \mathbf{x} + \frac{1}{O} \sum_{o=1}^O \left( \sum_{i=1}^R z_{o,i} \cdot \mathbf{b}_i \mathbf{a}_i^\top \right) \mathbf{x}, \quad (6)$$

where  $\Delta \mathbf{W}_o = \sum_{i=1}^R z_{o,i} \cdot \mathbf{b}_i \mathbf{a}_i^\top$  is the object-specific adapter for object  $o$ ,  $z_{o,i}$  is the  $i$ -th component of  $\mathbf{z}_o$ ,  $\mathbf{x}$  is the input feature, and  $\mathbf{W}_0$  is the frozen base weight. MoGA shares the same  $\{\mathbf{a}_i, \mathbf{b}_i\}$  across all objects. This Siamese structure enables efficient object-specific adaptation without parameter duplication.

**Integration and Training.** We integrate MoGA into SAM2’s memory attention module, specifically into the

linear projections for self-attention (Q, K, V) and cross-attention (Q). Each projection has its own selection of (rank-1) components, though the low-rank adapter and the gating module are shared across objects. Training uses a standard segmentation loss

$$\mathcal{L}_{\text{total}} = \frac{1}{T \cdot O} \sum_{t=1}^T \sum_{o=1}^O \mathcal{L}_{\text{seg}}(y_{o,t}, \hat{y}_{o,t}), \quad (7)$$

where  $\mathcal{L}_{\text{seg}}$  combines focal and dice losses,  $\hat{y}_{o,t}$  is the predicted mask for object  $o$  at frame  $t$  produced with its memory-conditioned adapter  $\Delta \mathbf{W}_o$ , and  $y_{o,t}$  is the respective ground-truth mask. Notably, the gating modules that determine which (rank-1) components to activate are not directly supervised. They learn to select appropriate adaptation paths solely through the segmentation loss, following established practices in mixture-of-experts literature [9].

## 5. Experiments

### 5.1. Implementation Details

We adopt the pre-trained SAM2 [37] and freeze all parameters except for the newly introduced MoGA modules and LayerNorm layers. MoGA modules are attached to the projection layers of the memory attention (key, query, and value of self-attention and query of cross-attention) and are trained jointly with LayerNorm, following [6, 36, 45, 48]. The parameters are optimized with AdamW [31] with a learning rate of  $5 \times 10^{-6}$ , weight decay of 0.1, batch size of 4, and a maximum of 3 objects per clip. The adapter rank  $R$  is set to 128. We use a temperature  $\tau$  of 0.3 for the Gumbel-sigmoid with a linear annealing to stabilize hard gating.

### 5.2. Results on Our Real-World Benchmark

We evaluate each model on our real-world benchmark, MVSeg-adv [16] and ACDC-Video [39], captured under real-world corruptions. Table 2 reports zero-shot segmentation results. The original SAM2, despite its strong performance on clear videos, shows substantial degradation under real corruptions, achieving only 69.6% and 63.5%  $\mathcal{J}\&\mathcal{F}$



Figure 4. Qualitative results on the real-world corrupted sequences of our benchmark. Each color indicates tracked objects: red for vehicles and green for pedestrians.

on MVSeg-adv and ACDC-Video, respectively. Restoration methods (URIE+SAM2, AirNet+SAM2) offer marginal improvements and even degrade performance. Frame-wise application of GaRA similarly achieves small improvements or degrades performance, reaching 69.7% and 61.3%  $\mathcal{J}\&\mathcal{F}$  on MVSeg-adv and ACDC-Video. By contrast, MoGA combined with SAM2 (MoGA+SAM2) achieves significant gains across all metrics, reaching 71.8%  $\mathcal{J}\&\mathcal{F}$  on MVSeg-adv and 64.5%  $\mathcal{J}\&\mathcal{F}$  on ACDC-Video. This demonstrates the effectiveness of memory-object-conditioned robustification over frame-level approaches. The improvements are consistent across both datasets, suggesting that MoGA generalizes well to diverse real-world corruption types without dataset-specific tuning.

### 5.3. Results on Synthetic Datasets

We additionally evaluate each model on synthetic corruption datasets, especially on the validation split of YouTube-VOS-C. Table 2 shows that SAM2 achieves 78.7%  $\mathcal{J}\&\mathcal{F}$  on YouTube-VOS-C, exhibiting a performance degradation from its clear performance, 82.2%. Frame-wise robustification methods, including restoration approaches and GaRA,

Table 3. Comparison with fully fine-tuned SAM2 on MVSeg-adv.

Method	Learnable Params	GPU Memory	$\mathcal{J}\&\mathcal{F}$
Fine-tuned SAM2	80.9M	25GB	71.5
MoGA+SAM2	1.1M	22GB	<b>71.8</b>

show performance deterioration. MoGA+SAM2 outperforms all baselines, achieving 79.9%  $\mathcal{J}\&\mathcal{F}$  on corrupted videos. On clean videos, MoGA+SAM2 reaches 82.6%, on par with the original SAM2 (82.2%). These results suggest that memory-object-based conditioning is more robust compared to frame-wise methods while maintaining reasonable performance on clear videos.

### 5.4. Qualitative Results

Figure 4 presents qualitative results for SAM2, URIE+SAM2, GaRA+SAM2, and MoGA+SAM2 on corrupted real-world videos from our benchmark. SAM2 predicts noisy masks or loses tracking of objects entirely.

Table 4. Comparison with LoRA.

Method	$\mathcal{J}\&\mathcal{F}$
SAM2 [37]	69.6
LoRA [14]+SAM2	70.9
MoGA+SAM2 (Ours)	<b>71.8</b>

Table 5. Results on short vs. long videos.

	SAM2	Ours
Short ( $\sim 6s$ )	69.6	<b>71.8</b>
Long ( $\sim 42s$ )	52.3	<b>56.2</b>




Figure 5. Visualization of gating masks over time. *Left*: input video frames from a nighttime driving sequence. *Right*: corresponding binary gating masks showing activated components (black) for tracked objects.

The image restoration method, URIE+SAM2, shows limited improvements and produces partial masks, capturing only fragments of objects under severe corruptions. Frame-wise application of GaRA, GaRA+SAM2, exhibits temporally inconsistent masks across frames due to its lack of temporal consistency. In contrast, MoGA+SAM2 shows robust predictions across corrupted frames. MoGA maintains temporal consistency by leveraging stored object memory, preventing inconsistent results.

### 5.5. Comparison with Full Fine-tuning

To evaluate parameter efficiency, we compare MoGA combined with SAM2 to a fully fine-tuned SAM2 in Table 3. The full fine-tuning baseline optimizes all 80.9M parameters of SAM2, achieving 71.5%  $\mathcal{J}\&\mathcal{F}$  on MVSeg-adv while requiring 25GB of GPU memory for training. In contrast, MoGA+SAM2 trains only 1.1M parameters yet achieves 71.8%  $\mathcal{J}\&\mathcal{F}$  using only 22GB of memory. This demonstrates that our memory-conditioned robustification is not only more parameter-efficient but also more effective than full fine-tuning of SAM2.

### 5.6. Comparison with LoRA.

We compare MoGA against LoRA [14] applied to SAM2 under the same training setup. As shown in Table 4, SAM2+LoRA achieves 70.9%  $\mathcal{J}\&\mathcal{F}$  on MVSeg-adv, im-




Figure 6. Quantitative and qualitative results during inference. *Top*:  $\mathcal{J}\&\mathcal{F}$  scores across real-world nighttime frames. *Bottom*: qualitative comparison of segmentation results in the sequence.

proving over the SAM2 baseline (69.6%) but underperforming MoGA+SAM2 (71.8%). LoRA updates weights uniformly, without any awareness of which object is being tracked or what the memory bank contains, applying identical adaptation. MoGA surpasses LoRA through its memory-object-conditioned gating, which enables temporally consistent and object-specific adaptation essential for RobustPVOS but absent in LoRA.

### 5.7. In-Depth Analysis

**Results on Long Videos.** To assess scalability to longer sequences, we construct extended videos by merging clips ( $\sim 6s$  each) from the same video into longer sequences ( $\sim 42s$ , 1K frames) on MVSeg-adv. As shown in the Table 5, SAM2 suffers a notable drop on longer videos. MoGA+SAM2, however, maintains improvement over SAM2 on both short and long video settings, suggesting that memory-conditioned gating remains stable as the memory bank grows.

**Impact of Gating Mask.** We investigate the effect of memory-conditioned gating by visualizing the learned gat-

Table 6. Effect of individual components of MoGA on performance on MVSeg-adv.

Memory-cond	Object-cond	$\mathcal{J}\&\mathcal{F}$
		69.6
✓		70.9
✓	✓	<b>71.8</b>

Table 7. Effect of adapter rank  $R$  of MoGA on performance on YouTube-VOS-C.

Metric	Rank $R$				
	32	64	<b>128</b>	256	512
$\mathcal{J}\&\mathcal{F}$	79.3	79.4	<b>79.9</b>	79.8	79.7

Table 8. Effect of temperature  $\tau$  of MoGA on YouTube-VOS-C.

Metric	Temperature $\tau$			
	0.1	0.3	0.5	0.7
$\mathcal{J}\&\mathcal{F}$	<b>79.9</b>	<b>79.9</b>	79.7	79.8

ing masks across frames in Figure 5. The visualization reveals that gating masks maintain temporal consistency across the video sequence. Despite this overall consistency, we observe small adaptive changes in the activated (rank-1) components over time. These smooth transitions in gating patterns demonstrate that our approach successfully leverages temporal information for stable robustification.

**Progressive Improvement of MoGA.** Figure 6 shows quantitative and qualitative comparisons between SAM2 and MoGA+SAM2 model on a sample sequence. While SAM2 performs around 40%  $\mathcal{J}\&\mathcal{F}$  across the sequence, MoGA progressively improves performance to over 80%  $\mathcal{J}\&\mathcal{F}$  by the end of the sequence. This behavior arises from the accumulation of object pointers in the memory bank, enabling the gating module to select progressively more suitable components of the adapter. The qualitative results also demonstrate that MoGA evolves from fragmented masks to complete segmentation, demonstrating the effectiveness of memory-conditioned progress during inference.

**Impact of Memory-object Conditioning.** Table 6 investigates the contribution of each component of MoGA on MVSeg-adv. The three variants differ as follows: the first uses no conditioning at all; the second applies memory-conditioning only, which aggregates all object pointers into a single shared representation before gating; the third is the full MoGA, which combines memory- and object-conditioning so that gating is performed independently for each object pointer. Without either form of conditioning, the model achieves only 69.6%  $\mathcal{J}\&\mathcal{F}$ . Incorporating temporal information via memory-conditioning improves performance to 70.9%  $\mathcal{J}\&\mathcal{F}$ , demonstrating the importance of object memory for maintaining cross-frame consistency. The full MoGA, which additionally conditions gating on each individual object pointer, achieves 71.8%  $\mathcal{J}\&\mathcal{F}$ , validating the effect of object-specific adaptation on top of temporal conditioning.

**Analysis on Rank of Adapter.** Table 7 shows the effect of adapter rank  $R$  on YouTube-VOS-C. The performance remains stable across different rank values, with our default setting of  $R = 128$  achieving optimal performance at 79.9%

$\mathcal{J}\&\mathcal{F}$ . Lower ranks ( $R = 32, 64$ ) yield reasonable performance, showing that compact representations can capture object-specific patterns. Increasing the rank to 256 or 512 yields a degradation, indicating saturation beyond  $R = 128$ . The optimal rank of 128 strikes a balance between adaptation flexibility and parameter efficiency, introducing only 1.1M trainable parameters while achieving substantial improvements in robustness.

**Analysis on Temperature of Gumbel-Sigmoid.** Table 8 shows the impact of temperature  $\tau$  of Gumbel-Sigmoid in Eq. (3). Lower values of  $\tau$  produce harder component selection, while higher values yield softer activations. We find that MoGA is insensitive to  $\tau$ , with 0.1 and 0.3 yielding the best performance on YouTube-VOS-C. Thus, our memory-based conditioning provides strong guidance for gating decisions, enabling hard selection without training instability.

## 6. Conclusion

We have presented the first comprehensive study on robust promptable video object segmentation, establishing it as a central research direction for real-world video object segmentation. Our first contribution is a carefully designed benchmark comprising real-world evaluation datasets with dense object-level annotations under real adverse conditions and synthetic datasets with controlled temporal degradations. We have further proposed Memory-object-conditioned Gated-rank Adaptation (MoGA) as an effective baseline that leverages object representations from the memory bank to achieve object-specific robustification while maintaining temporal consistency across frames. Extensive experiments demonstrate that memory-based conditioning substantially outperforms per-frame approaches, validating its effectiveness for handling spatially and temporally varying degradations in videos. Our comprehensive benchmark and baseline provide foundations for future research in RobustPVOS.

**Acknowledgement.** This work was supported by AI Graduate School Program at POSTECH (RS-2019-II191906), and the IITP grants (RS-2024-00457882, RS-2022-II220926, RS-2024-00509258, RS-2024-00469482) funded by the Ministry of Science and ICT, Korea.

## References

- [1] Yuang Ai, Huaibo Huang, and Ran He. Lora-ir: Taming low-rank experts for efficient all-in-one image restoration. *arXiv preprint arXiv:2410.15385*, 2024. 2
- [2] Anurag Bagchi, Zhipeng Bao, Yu-Xiong Wang, Pavel Tokmakov, and Martial Hebert. Refereverything: Towards segmenting everything we can speak of in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23221–23231, 2025. 2
- [3] Qi Bi, Shaodi You, and Theo Gevers. Generalized foggy-scene semantic segmentation by frequency decoupling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1389–1399, 2024. 2
- [4] Wei-Ting Chen, Yu-Jiet Vong, Sy-Yen Kuo, Sizhou Ma, and Jian Wang. RobustSAM: segment anything robustly on degraded images. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3
- [5] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3151–3161, 2024. 2
- [6] Thomas De Min, Massimiliano Mancini, Karteek Alahari, Xavier Alameda-Pineda, and Elisa Ricci. On the effectiveness of layernorm tuning for continual learning in vision transformers. In *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 3577–3586. IEEE, 2023. 5
- [7] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, Philip HS Torr, and Song Bai. MOSE: A new dataset for video object segmentation in complex scenes. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 3
- [8] Shuangrui Ding, Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Yuwei Guo, Dahua Lin, and Jiaqi Wang. Sam2long: Enhancing sam 2 for long video segmentation with a training-free memory tree. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 1
- [9] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022. 5
- [10] Volodymyr Fedynyak, Yaroslav Romanus, Bohdan Hlovatskyi, Bohdan Sydor, Oles Dobosevych, Igor Babin, and Roman Riazantsev. Devos: Flow-guided deformable transformer for video object segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 240–249, 2024. 2
- [11] Diandian Guo, Deng-Ping Fan, Tongyu Lu, Christos Sakaridis, and Luc Van Gool. Vanishing-point-guided video semantic segmentation of driving scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3544–3553, 2024. 2
- [12] Pinxue Guo, Wanyun Li, Hao Huang, Lingyi Hong, Xinyu Zhou, Zhaoyu Chen, Jinglun Li, Kaixun Jiang, Wei Zhang, and Wenqiang Zhang. X-prompt: Multi-modal visual prompt for video object segmentation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 5151–5160, 2024. 2
- [13] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *Proc. International Conference on Learning Representations (ICLR)*, 2019. 1, 3
- [14] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *Proc. International Conference on Learning Representations (ICLR)*, 2022. 4, 7
- [15] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proc. International Conference on Learning Representations (ICLR)*, 2017. 5
- [16] Wei Ji, Jingjing Li, Cheng Bian, Zongwei Zhou, Jiaying Zhao, Alan L. Yuille, and Li Cheng. Multispectral video semantic segmentation: A benchmark dataset and baseline. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 3, 5
- [17] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2023. 2
- [18] Taewoo Kim, Jeongmin Lee, Lin Wang, and Kuk-Jin Yoon. Event-guided deblurring of unknown exposure time videos. In *European Conference on Computer Vision*, pages 519–538. Springer, 2022. 3
- [19] Taeh Kim, Jinhyung Kim, Minho Shim, Sangdoo Yun, Myunggu Kang, Dongyoon Wee, and Sangyoun Lee. Exploring temporally dynamic data augmentation for video recognition. In *Proc. International Conference on Learning Representations (ICLR)*, 2023. 1, 3
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2
- [21] Sohyun Lee, Taeyoung Son, and Suha Kwak. Fifo: Learning fog-invariant features for foggy scene segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [22] Sohyun Lee, Jaesung Rim, Boseung Jeong, Geonu Kim, Byungju Woo, Haechan Lee, Sunghyun Cho, and Suha Kwak. Human pose estimation in extremely low-light conditions. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [23] Sohyun Lee, Namyup Kim, Sungyeon Kim, and Suha Kwak. Frest: Feature restoration for semantic segmentation under multiple adverse conditions. In *Proc. European Conference on Computer Vision (ECCV)*. Springer, 2024.
- [24] Sohyun Lee, Yeho Gwon, Lukas Hoyer, and Suha Kwak. GaRA-SAM: Robustifying segment anything model with gated-rank adaptation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2025. 2, 4, 5
- [25] Sohyun Lee, Nayeong Kim, Juwon Kang, Seong Joon Oh, and Suha Kwak. Testdg: Test-time domain general-

- ization for continual test-time adaptation. *arXiv preprint arXiv:2504.04981*, 2025. 2
- [26] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-In-One Image Restoration for Unknown Corruption. In *IEEE Conference on Computer Vision and Pattern Recognition*, New Orleans, LA, 2022. 2, 5
- [27] Hebei Li, Jin Wang, Jiahui Yuan, Yue Li, Wenming Weng, Yansong Peng, Yueyi Zhang, Zhiwei Xiong, and Xiaoyan Sun. Event-assisted low-light video object segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [28] Minghan Li, Shuai Li, Xindong Zhang, and Lei Zhang. UniVS: Unified and universal video segmentation with prompts as queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3227–3238, 2024. 1, 2
- [29] Xin Li, Deshui Miao, Zhenyu He, Yaowei Wang, Huchuan Lu, and Ming-Hsuan Yang. Learning spatial-semantic features for robust video object segmentation. In *Proc. International Conference on Learning Representations (ICLR)*, 2024. 2
- [30] Yang Liu, Yufei Yin, Chenchen Jing, Muzhi Zhu, Hao Chen, Yuling Xi, Bo Feng, Hao Wang, Shiyu Li, and Chunhua Shen. Unified open-world segmentation with multi-modal prompts. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 2
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. International Conference on Learning Representations (ICLR)*, 2019. 5
- [32] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7086–7096, 2022. 2
- [33] Haiyang Mei, Pengyu Zhang, and Mike Zheng Shou. Sam-i2v: Upgrading sam to support promptable video segmentation with less than 0.2% training cost. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 3417–3426, 2025. 1, 2, 4
- [34] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 3
- [35] Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. PromptIR: Prompting for all-in-one image restoration. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2023. 2
- [36] Wang Qi, Yu-Ping Ruan, Yuan Zuo, and Taihao Li. Parameter-efficient tuning on layer normalization for pre-trained language models. *arXiv preprint arXiv:2211.08682*, 2022. 5
- [37] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Juntao Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollar, and Christoph Feichtenhofer. SAM 2: Segment anything in images and videos. In *Proc. International Conference on Learning Representations (ICLR)*, 2025. 1, 2, 4, 5, 7
- [38] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1
- [39] Christos Sakaridis, Haoran Wang, Ke Li, René Zurbrugg, Arpit Jadon, Wim Abbeloos, Daniel Olmeda Reino, Luc Van Gool, and Dengxin Dai. ACDC: The adverse conditions dataset with correspondences for robust semantic driving scene perception. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 1, 3, 5
- [40] Hongje Seong, Junhyuk Hyun, and Euntai Kim. Kernelized memory network for video object segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 2
- [41] Taeyoung Son, Juwon Kang, Namyup Kim, Sunghyun Cho, and Suha Kwak. Urie: Universal image enhancement for visual recognition in the wild. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 2, 5
- [42] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning video object segmentation with visual memory. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2
- [43] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning motion patterns in videos. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [44] Yi-Hsuan Tsai, Ming-Hsuan Yang, and Michael J Black. Video segmentation via object flow. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [45] Taha ValizadehAslani and Hualou Liang. Layernorm: A key component in parameter-efficient fine-tuning. *arXiv preprint arXiv:2403.20284*, 2024. 5
- [46] Yunyang Xiong, Chong Zhou, Xiaoyu Xiang, Lemeng Wu, Chenchen Zhu, Zechun Liu, Saksham Suri, Balakrishnan Varadarajan, Ramya Akula, Forrest Iandola, et al. Efficient track anything. In *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 1
- [47] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. 1, 3, 4, 5
- [48] Bingchen Zhao, Haoqin Tu, Chen Wei, Jieru Mei, and Cihang Xie. Tuning layernorm in attention: Towards efficient multi-modal llm finetuning. *arXiv preprint arXiv:2312.11420*, 2023. 5
- [49] Junbao Zhou, Ziqi Pang, and Yu-Xiong Wang. Rmem: Restricted memory banks improve video object segmentation. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 4
- [50] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2023. 2